arXiv:1601.02713v1 [physics.chem-ph] 12 Jan 2016

On the hierarchical parallelization of *ab initio* simulations

Sergi Ruiz-Barragan[a,b], Kazuya Ishimura[b], Motoyuki Shiga[a,*]

[a] *CCSE, Japan Atomic Energy Agency, 178-4-4, Wakashiba, Kashiwa, Chiba, 277-0871, Japan*

[b] *Department of Theoretical and Computational Molecular Science, Institute for Molecular Science, Myodaiji, Okazaki, Aichi 444-8585, Japan*

Abstract

A hierarchical parallelization has been implemented in a new unified code PIMD-SMASH for *ab initio* simulation where the replicas and the Born-Oppenheimer forces are parallelized. It is demonstrated that *ab initio* path integral molecular dynamics simulations can be carried out very efficiently for systems up to a few tens of water molecules. The code was then used to study a Diels-Alder reaction of cyclopentadiene and butenone by *ab initio* string method. A reduction in the reaction energy barrier is found in the presence of hydrogen-bonded water, in accordance with experiment.

Keywords: *ab initio* simulation, hierarchical parallelization, SMASH, PIMD, string method, Diels-Alder reaction, butenone, cyclopentadiene, hydrogen bond

1. Introduction

In recent years, the massively parallel architecture has played a key role in the advances of modern supercomputers, such as the K computer [1]. In computational chemistry, much effort has been devoted toward the development of theoretical methods and numerical algorithms to make best use of their parallel performance. Among them, there is a class of molecular simulations where the

*Corresponding author

Email address: shiga.motoyuki@jaea.go.jp (Motoyuki Shiga)

Born-Oppenheimer (BO) energies and the BO forces can be evaluated in parallel for a set of molecular structures of the same system, which are referred to as the "replicas", "beads", or "images" depending on the method used. This class of molecular simulations is employed for many purposes, such as quantum statistics (e.g. path integral molecular dynamics (path integral MD) [2]), statistical sampling (e.g. replica exchange MD [3]), nonadiabatic dynamics (e.g. surface hopping dynamics [4]), reaction paths (e.g. Onsager-Machlup method [5], string method [6, 7]) or free energy exploration methods (e.g. constrained MD [8], multiple walker metadynamics [9, 10]).

Although many of the above methods have been established for empirical force fields, they are expected to be replaced by BO forces based on *ab initio* or first-principles electronic structure calculations, for accurate theoretical prediction. For this aim, a fortran-based interface code called PIMD[11] was originally developed in our research group for *ab initio* path integral simulations [12], see reference [13] for a short explanation of the code. The latest version of PIMD supports other replica simlations such as the string method, multiple-walker metadynamics or replica exchange methods. The PIMD code has a subroutine to extract the BO forces from other softwares, such as Turbomole[14], Mopac, Gaussian[15, 13], etc., using the intrinsic subroutine 'system' that passes unix commands to the shell and reading the resulting data from the hard disk. For the similar purpose, a python-based interface i-Pi code[16] was released to carry out path integral simulations extracting BO forces from FHI-aims [17] and CP2K [18]. For an efficient *ab initio* simulations with many replicas, hierarchical parallelization in the replicas ("replica-parallel") and in the BO forces ("force-parallel") should be required. Actually, this function has been employed in CPMD code based on the plane-wave density functional theory [19]. However, calling a set of independent parallel force calculations via 'system' subroutine as done in the PIMD code is often prohibited in public facilities, such as the K computer, for administrative reasons. Therefore, to enable such a hierarchical parallelization, we have made a coherent unification of the PIMD code and the *ab initio* electronic structure code, Scalable Molecular Analysis Solver for High-

performance computing systems (SMASH) [20]. The new PIMD-SMASH code is a single MPI program for *ab initio* replica simulation based on conventional quantum chemistry with Gaussian basis sets. In this paper, we illustrate the technical details on how to integrate parallel molecular simulation and *ab initio* electronic structure codes that have already been parallelized. We study the advantages and limitations on the parallel performance of hierarchical parallelized code for two different simulations, i.e. path integral MD simulation of water clusters and minimum energy path (MEP) calculation of Diels-Alder (DA) reaction of cyclopentadiene and butenone using the string method. Finally, scientific discussions have been made on the results of the DA reaction which is a typical example of green organic synthesis in an aqueous environment.

2. Methods and implementation

The code of SMASH was adapted and introduced in PIMD code to obtain a single program, the PIMD-SMASH. The PIMD-SMASH code is schematically shown in Figure 1(a). In PIMD a set of n nuclear configurations (replicas) is generated. Then, in SMASH the BO energies and forces of the respective replica are computed using m cores. Thus, the hierarchically parallel computation is done using $m \times n$ cores in total. In PIMD the BO energies and forces are gathered, all the positions of n replicas are updated according to the simulation method chosen. Finally, all the replicas are used to start a next step of the MD cycle.

The PIMD-SMASH code uses the MPI library routines, see the pseudocode shown in Figure 1(b). The hierarchical parallelization requires a group of replica-parallel, mpi_group_PIMD and a group of force-parallel, mpi_group_SMASH, as shown in Figure 1(c). Accordingly, the two respective MPI communicators, mpi_comm_PIMD and mpi_comm_SMASH, are created at the beginning of the calculation. The process ID's for the respective communicators, $myrank_PIMD$ and $myrank_SMASH$, are generated to distribute computational tasks; $myrank_PIMD$ is ranged from 0 to $n-1$ while $myrank_SMASH$

is ranged from 0 to $m-1$. The BO energies and forces of all the replicas are shared among all the processes. This is done with the simple sum using the $MPI_ALLREDUCE$ routine, since the data is zero wherever not computed. Using the BO energies and forces, the positions and velocities are updated. The updates are not parallelized thus the improvement in the efficiency will be insignificant because the bottleneck is the force calculation. However, future updates in the code will go in this direction. To share the information of input files first read by the parent process, MPI_BCAST communication is required only once at the beginning of the calculation. Although SMASH have an initial implementation of hybrid parallelization MPI/OpenMP, this version of PIMD-SMASH does not have this feature available. This option will be introduced in the next version of PIMD-SMASH using the improvement of the next version of SMASH.

Thus, PIMD-SMASH is a seamless MPI code having the electronic structure calculations included, as opposed to the original version of the PIMD code. This not only makes the code more efficient but also more friendly to both users and administrators (MPI packages have readily accessible documentation and easy tools to manage parallel programs using a job scheduler). Furthermore, it is easier to extend the functions of the code, and to optimize the efficiency. For instance, the molecular orbitals of each replica saved in the memory can be reused for the initial guess of the next MD step, without hard disk access.
We will publish the source code freely with GNU General Public License version 3 [21] in 2016, and a web page is under construction in Japan Atomic Energy Agency (JAEA). Further details of implementation will be available with the source code. The PIMD code facilitates the implementation of user-made forces. In fact, PIMD could also be readily coupled to other electronic structure codes to take advantage of supported sampling methods.

3. Results and discussion

We have completed the coding of PIMD-SMASH in the way described in the previous section. In this section, the parallel performance is tested to clarify its advantages and limitations on a quantitative basis. All the calculations were performed on the Fujitsu PRIMERGY BX900 machine composed of Intel Xeon X5570 CPUs (8 cores/node) in JAEA.

3.1. Parallel performance: path integral MD of water clusters

Short runs of *ab initio* path integral MD simulation [22] of 1–64 water molecules (in the linearly aligned configurations) were carried out to study the parallel performance. The electronic structure calculations were based on the B3LYP/6-31G* level of theory, where 18 basis functions are allocated for each water molecule. Following the default settings of SMASH, the self-consistent field (SCF) cycles with the density threshold of 1.0×10^{-6} and the integral cutoff of 1.0×10^{-12} were used.

The parallel efficiency with respect to the beads (replicas, in path integral MD simulations) is shown in Table 1 and 2. (Note that in the conventional path integral MD simulations of water molecules at the room temperature, it is known that the number of beads required is about 16–64 to converge the average energies and distribution functions [23], although there are recent efforts to reduce the number of beads [24, 25].) Table 1 shows the average time per step when the number of cores is increased in proportion to the number of beads, i.e. the weak scaling (scaling with fixed system size per processor), while Table 2 shows the performance for 64 beads with various number of cores which shows the strong scaling (scaling at fixed system size). In both cases, four cores were used for the force calculations such that the results are unaffected by the force-parallel efficiency. A small but gradual increase of the average time shown in Table 1 indicates that the replica-parallel efficiency is high but not perfect. After extensive tests, it turned out that this deterioration stems from the communication overflow when many nodes were occupied by an independent set of

SMASH calculations. This problem is heavily machine dependent, maybe particular to BX900, since we find the replica-parallel efficiency was nearly perfect when the same test was done on other computers. Even so, the difference of about three seconds between the cases of two beads and 64 beads is not very serious. On the other hand, the strong scaling is above 90% for the cases shown in Table 2.

Next, the parallel performance of *ab initio* path integral MD simulations were tested in the case of two beads. In Figure 2(a), we plot the results of the relative speed, which is defined as the steps proceeded per unit time in comparison with those using a single core for the same system. One can see a general trend that the parallel performance is better as the system size increases. For instance, the relative speed for a single water molecule stays around 15 when 32 or more cores are used. This is because the amount of calculation is too small to be parallelized efficiently. On the other hand, the relative speed of 16 water molecules grows from 28.7 to 91.1, which amounts to the parallel efficiency as high as 71.2%.

To understand the origin of the inefficiency, the relative speed is separated between the contributions from the BO energies and the BO forces, which are shown in Figure 2(b) and 2(c), respectively. The force calculation clearly outperforms the energy calculation. The parallel efficiencies in the BO force calculations were found to be over 70% in all the cases studied. In contrast, the efficiencies are poorer in the BO energy calculations, with less than 50% beyond 256 cores. To see how this feature affects the total timing, the ratio of calculation time spent in the BO force and the BO energy is displayed in Figure 2d. Here one finds that the ratio is large for a system of more than 32 water molecules, especially in the case of small number of cores, 1–64. This explains the fact that the efficiency remains high up to 64 cores for large systems. However, the ratio drops quickly beyond 64 cores due to the inefficiency of the energy calculation. Thus, the total parallel efficiency decreases as the average time per step is dominated by the BO energy contribution.

In Figure 3, we show the average time per step in path integral MD simula-

tion. This index is crucial for any kind of MD simulations. For instance, it is often the case that the (path integral) MD trajectory of several tens of picoseconds, which amounts to a few hundred thousands of steps, is required to obtain statistically converged results of the pair distributions of liquid water. Say, to complete 100000 steps in 30 days of continuous run, the average time per step must be less than 25.92 seconds. This value is marked as a dotted line in Figure 3. One can see that the results are below the line for a single core for the small systems up to 4 water molecules. For 8, 16 and 32 water molecules, there is a crossover at around 4, 19 and 210 cores, respectively. In this way, as the system size increases the number of cores required is 'delayed' to a larger value. Finally, in the case of 64 water molecules, the crossover is not found for any number of cores, which means that such a calculation can never be efficiently executed. This is the reason why the MD simulation of large systems remains as a challenging issue.

3.2. Parallel performance: string method for DA reaction

Next, hierarchical parallelization is applied to the string method for a simulation of the DA reaction of cyclopentadiene and butenone. The string is described in terms of 120 images (replicas in the string method) that interconnect the reactant minimum to the product minimum, which are shown later in Figure 5. Here, the electronic structure calculations are based on the B3LYP/6-31G** level of theory. The results are summarized in Figure 4. Using a single core it took 38159 seconds per step, where a step in the string method refers to the force calculation of all the images followed by the update of the whole string. With 120 cores, the average times per step reduced to 244.8–318.0 seconds depending on how the cores are distributed. In hierarchical parallelization, the cores can be used in different products of the force-parallel number m and the image-parallel number n, and the best combination seems to be moderate values, $n = 15$ and $m = 8$ in the present case. Clearly, the image-parallel has better load balancing and it is advantageous in reducing the communication during the force calculations, but it is also disadvantageous in calculating many

images at the same time and waiting for all to finish. This aspect is not negligible when the number of SCF cycles varies among the images, which is the case between equilibrium and transition state. This point is particularly important in the string method because the images are usually in very different geometries, but it may become less serious in other methods such as path integral MD simulations. Comparing the fully image-parallel cases, the time per step of the double force-parallel using 240 cores, was almost a half of that of the single force-parallel, using 120 cores. This confirms that the efficiency of the string method when the force-parallel number is sufficiently small.

3.3. Application: Diels-Alder reaction

Finally, we have analyzed the MEP of the DA reaction of cyclopentadiene and butenone obtained by the string method using hierarchical parallelization. This reaction is known as one of the earliest organic reactions found to be accelerated in aqueous solution [26]. Initially, it was proposed that the acceleration is due to hydrophobicity [26]. More recently, some theoretical works [27, 28, 29, 30] suggested that the hydrogen bond plays the key role in the acceleration. All the MD simulations of aqueous solutions done so far are based on the semi-empirical methods [28, 29, 30] since the statistical accuracy less than 1 kcal/mol would be formidable for *ab initio* MD. To our knowledge, the only *ab initio* study is reference [28] where Hartree-Fock calculations of the molecular cluster system with one water molecule were completed using the constrained geometry optimization. Here we present the *ab initio* calculation of the molecular cluster system with different numbers water molecules (0-2). For each system, the full MEP from the reactant minimum to the product minimum is obtained without any geometrical constraints, which allows the estimation of the activation energies of forward and backward reactions. To account for the electron correlation effects, the calculation is based on the B3LYP/6-31G** level of theory.

In Figure 5, the potential energy profiles along with the structures of stationary states are shown. We find that the activation energy of the forward reaction without water molecule is 16.9 kcal/mol, but it is lowered by 0.9 kcal/mol when

water molecule is added. There are two possible conformations when adding more water molecules, i.e. the doubly hydrogen bonded conformation or the hydrogen bonded chain conformation (see Figure 5). According to our calculation, the latter is more stable. It is favoured for only one water molecule to be directly hydrogen bonded to butenone. Moreover, the hydrogen bonded chain conformation has a lower activation energy than the doubly hydrogen bonded conformation. However, the activation energy of the cluster system with two water molecules in the hydrogen bonded chain conformation is lower by 0.8 kcal/mol compared to that without water, almost the same that one water molecule cluster. The reduction of activation energy are qualitatively consistent with the experimental finding that the DA reaction is accelerated in water. It is clear that the hydrogen bond plays the important part in the acceleration mechanism. However, small increment of water molecules does not modify the acceleration rate, and these values are smaller than the experimental estimate of the free energy barrier in liquid water [31, 32], where a change of 4 kcal/mol is found. Therefore, for more quantitative analysis, *ab initio* studies with higher accuracy, and inclusion of solvent effects either by implicit or explicit treatment of water molecules, etc. should be useful for a future work.

4. Conclusions

Unification of PIMD and SMASH codes has enabled the hierarchical parallelization of *ab initio* simulations with respect to replicas and forces. As mentioned in the Introduction, this technique is open to a broad class of *ab initio* simulations, such as quantum statistics, statistical sampling, nonadiabatic dynamics, reaction paths or free energy exploration methods. The parallel performance has been tested on the JAEA supercomputer with 2048 cores for two cases, i.e. path integral MD simulations of water molecules with various system sizes and reaction path calculation of the cyclopentadiene-butenone DA reaction by the string method. The key measure in the feasibility of these *ab initio* simulations is the average time per step. This index drops as the number of

cores are larger, and there is a minimum value of this measure reachable for a given system size. More critically, there is an upper limit in the system size for a required value of time per step. In the present case, this limitation comes from the deterioration of parallelization, in particular, the relative speed of the BO energy calculation shown in Figure 2b. The strong scaling of large systems should be addressed in future development of *ab initio* simulations. The communication between replica-parallel is not very important in the case of this study. In the hierarchical parallelization, there is a room of optimal choice of the replica- and force-parallel numbers. The data obtained from this work should be useful not only for the application studies using multiple replica simulation but also for a further development of parallel codes and algorithms to be expected in the future. Some improvements in the code are planning in a near future, such as the introduction of the hybrid parallelization MPI/OpenMP or the proper parallelization of coordination and velocities updating, to improve the shown efficiency.

The hierarchically parallelized code has been used to study the DA reaction of cyclopentadiene and butenone by the *ab initio* string method. The MEP is calculated in the absence of water, and in the presence of one and two water molecules hydrogen bonded to butenone. The results indicated a lowering of the energy barrier in the presence of water qualitatively consistent with the experimental finding that the reaction rate is higher in liquid water than in the organic solvents. For a more quantitative analysis including hydrophobic effects, *ab initio* calculations with improved accuracy and taking account of the solvent effects should be indispensable.

Acknowledgement

This work was supported by Theoretical and Computational Chemistry Initiative (TCCI) and Computational Materials Science Initiative (CMSI) in the Strategic Programs for Innovative Research, MEXT, Japan (Project ID: hp150231), and JSPS KAKENHI Grant Number 25730079. We thank the pri-

ority use of JAEA BX900 supercomputer in Center for Computational Science and E-systems at Tokai Research Establishment, JAEA.
We are grateful to Dr. A. Malins in JAEA for proofreading the manuscript of this paper.

References

References

[1] K computer (Fujitsu SPARC64 VIIIfx processor (2.0GHz, 8 cores/node). URL `http://www.aics.riken.jp/en/k-computer/about/`

[2] M. E. Tuckerman, B. J. Berne, G. J. Martyna, M. L. Klein, Efficient molecular dynamics and hybrid Monte Carlo algorithms for path integrals, The Journal of Chemical Physics 99 (4) (1993) 2796–2808.

[3] Y. Sugita, Y. Okamoto, Replica-exchange molecular dynamics method for protein folding, Chemical physics letters 314 (1) (1999) 141–151.

[4] J. C. Tully, Molecular dynamics with electronic transitions, The Journal of Chemical Physics 93 (2) (1990) 1061–1071.

[5] H. Fujisaki, M. Shiga, A. Kidera, Onsager–Machlup action-based path sampling and its combination with replica exchange for diffusive and multiple pathways, The Journal of chemical physics 132 (13) (2010) 134101.

[6] W. E, W. Ren, E. Vanden-Eijnden, String method for the study of rare events, Physical Review B 66 (2002) 052301.

[7] E. Weinan, W. Ren, E. Vanden-Eijnden, Simplified and improved string method for computing the minimum energy paths in barrier-crossing events, The Journal of Chemical Physics 126 (16) (2007) 164103.

[8] E. Carter, G. Ciccotti, J. T. Hynes, R. Kapral, Constrained reaction coordinate dynamics for the simulation of rare events, Chemical Physics Letters 156 (5) (1989) 472–477.

[9] A. Laio, M. Parrinello, Escaping free-energy minima, Proceedings of the National Academy of Sciences 99 (20) (2002) 12562–12566.

[10] P. Raiteri, A. Laio, F. L. Gervasio, C. Micheletti, M. Parrinello, Efficient reconstruction of complex free energy landscapes by multiple walkers metadynamics, The Journal of Physical Chemistry B 110 (8) (2006) 3533–3539.

[11] M. Shiga, PIMD version 1.7 (2015).

[12] M. Shiga, M. Tachikawa, S. Miura, Ab initio molecular orbital calculation considering the quantum mechanical effect of nuclei by path integral molecular dynamics, Chemical Physics Letters 332 (3) (2000) 396–402.

[13] M. Shiga, A. Nakayama, Ab initio path integral ring polymer molecular dynamics: Vibrational spectra of molecules, Chemical Physics Letters 451 (4) (2008) 175–181.

[14] TURBOMOLE V6.2 2010, a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH, 1989-2007, TURBOMOLE GmbH, since 2007; available from
`http://www.turbomole.com`.

[15] M. J. Frisch, G. W. Trucks, H. B. Schlegel, G. E. Scuseria, M. A. Robb, J. R. Cheeseman, G. Scalmani, V. Barone, B. Mennucci, G. A. Petersson, H. Nakatsuji, M. Caricato, X. Li, H. P. Hratchian, A. F. Izmaylov, J. Bloino, G. Zheng, J. L. Sonnenberg, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, J. A. Montgomery, Jr., J. E. Peralta, F. Ogliaro, M. Bearpark, J. J. Heyd, E. Brothers, K. N. Kudin, V. N. Staroverov, R. Kobayashi, J. Normand, K. Raghavachari, A. Rendell, J. C. Burant, S. S. Iyengar, J. Tomasi, M. Cossi, N. Rega, J. M. Millam, M. Klene, J. E. Knox, J. B. Cross, V. Bakken, C. Adamo, J. Jaramillo, R. Gomperts, R. E. Stratmann, O. Yazyev, A. J. Austin, R. Cammi, C. Pomelli, J. W. Ochterski, R. L. Martin, K. Morokuma, V. G. Zakrzewski, G. A. Voth,

P. Salvador, J. J. Dannenberg, S. Dapprich, A. D. Daniels, . Farkas, J. B. Foresman, J. V. Ortiz, J. Cioslowski, D. J. Fox, Gaussian 09 Revision D.01, gaussian Inc. Wallingford CT 2009.

[16] M. Ceriotti, J. More, D. E. Manolopoulos, i-PI: A Python interface for ab initio path integral molecular dynamics simulations, Computer Physics Communications 185 (3) (2014) 1019 – 1026.

[17] V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, M. Scheffler, Ab initio molecular simulations with numeric atom-centered orbitals, Computer Physics Communications 180 (11) (2009) 2175–2196.

[18] J. Hutter, M. Iannuzzi, F. Schiffmann, J. VandeVondele, CP2K: atomistic simulations of condensed matter systems, Wiley Interdisciplinary Reviews: Computational Molecular Science 4 (1) (2014) 15–25.

[19] W. Andreoni, A. Curioni, New advances in chemistry and material science with CPMD and parallel computing, Parallel Computing 26 (2000) 819–842.

[20] K. Ishimura, Scalable molecular analysis solver for high-performance computing systems (SMASH) (2014-2015).
URL `http://smash-qc.sourceforge.net/`

[21] Gnu general public license.
URL `http://www.gnu.org/licenses/gpl.html`

[22] M. Shiga, M. Tachikawa, S. Miura, A unified scheme for ab initio molecular orbital theory and path integral molecular dynamics, The Journal of Chemical Physics 115 (20) (2001) 9149–9159.

[23] W. Shinoda, M. Shiga, Quantum simulation of the heat capacity of water, Physical Review E 71 (4) (2005) 041204.

[24] K. Suzuki, M. Tachikawa, M. Shiga, Efficient ab initio path integral hybrid monte carlo based on the fourth-order trotter expansion: application to

fluoride ion-water cluster, The Journal of Chemical Physics 132 (14) (2010) 144108.

[25] M. Ceriotti, D. E. Manolopoulos, M. Parrinello, Accelerating the convergence of path integral dynamics with a generalized langevin equation, The Journal of chemical physics 134 (8) (2011) 084104.

[26] R. Breslow, Hydrophobic effects on simple organic reactions in water, Accounts of Chemical Research 24 (6) (1991) 159–164.

[27] J. F. Blake, D. Lim, W. L. Jorgensen, Enhanced hydrogen bonding of water to Diels-Alder transition states. ab initio evidence, The Journal of Organic Chemistry 59 (4) (1994) 803–805.

[28] J. F. Blake, W. L. Jorgensen, Solvent effects on a Diels-Alder reaction from computer simulations, Journal of the American Chemical Society 113 (19) (1991) 7430–7432.

[29] J. Chandrasekhar, S. Shariffskul, W. L. Jorgensen, QM/MM simulations for Diels-Alder reactions in water: Contribution of enhanced hydrogen bonding at the transition state to the solvent effect, The Journal of Physical Chemistry B 106 (33) (2002) 8078–8085.

[30] T. R. Furlani, J. Gao, Hydrophobic and hydrogen-bonding effects on the rate of Diels-Alder reactions in aqueous solution, The Journal of Organic Chemistry 61 (16) (1996) 5492–5497.

[31] D. C. Rideout, R. Breslow, Hydrophobic acceleration of Diels-Alder reactions, Journal of the American Chemical Society 102 (26) (1980) 7816–7817.

[32] J. B. Engberts, Diels-Alder reactions in water: Enforced hydrophobic interaction and hydrogen bonding, Pure and Applied Chemistry 67 (5) (1995) 823–828.

[33] O. Acevedo, W. L. Jorgensen, Understanding rate accelerations for Diels-Alder reactions in solution using enhanced QM/MM methodology, Journal of Chemical Theory and Computation 3 (4) (2007) 1412–1419.

Table 1: Average time per step in the path integral MD simulation of eight water molecules with various number of beads. Four cores are used per bead. All the beads are calculated in parallel.

Number of cores	8	16	32	64	128	256
Number of beads	2	4	8	16	32	64
Time average [s]	20.6	20.7	21.4	22.1	22.8	23.9

Table 2: Relative speed in the path integral MD simulation of eight water molecules using various number of cores. Four cores are used per bead.

Number of cores	8	16	32	64	128	256
Number of beads	64	64	64	64	64	64
Relative speed	1	2.0	4.0	7.8	15.2	28.87

Table 3: Transition state stabilization energy in water with respect to the gas phase reaction for DA reaction of cyclopentadiene and butenone of different works. Energies in kcal/mol. Free energies at 25°C and 1 atm.

Method	$\Delta\Delta U^{\ddagger}$	$\Delta\Delta G^{\ddagger}$	Reference
B3LYP/6-31G**	-0.9[a]		This work
B3LYP/6-31G**	-0.8[b]		This work
HF/6-31G*	-2.23[a]		Ref[27]
PDDG/PM3/MM/MC		-3.5[c]	Ref[33]
AM1/TIP3P		-3.5[c]	Ref[30]
AM1-OPLS-CM1A		-2.8[c]	Ref[29]
Experimental		-3.8[d]	Ref[31]
Experimental		-2.4[e]	Ref[32]

[a] With 1 water molecule.

[b] With 2 water molecules.

[c] Solvated in explicit water.

[d] Relative to isooctane at 20°C.

[e] Relative to 1-propanol.

Figure 1: a) Schematic of the hierarchical parallelized code PIMD-SMASH. b) Pseudo-code of the unified program PIMD-SMASH in Fortran using MPI subroutines. The comments are marked with "!". c) Scheme of the MPI processes (numbers), the MPI groups and the MPI communications (coloured rectangle). The red colour indicates the PIMD communication while the blue indicates the SMASH communications of the process '0'.

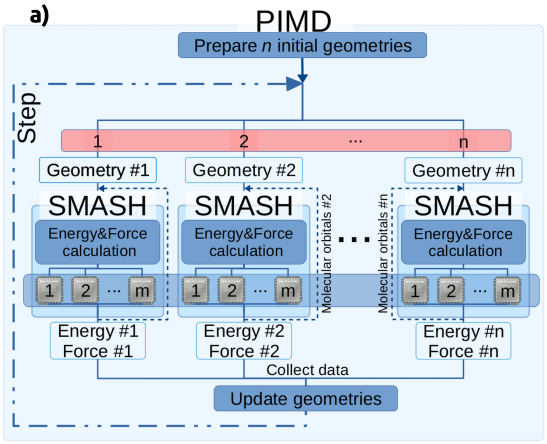

b)

```
program PIMD-SMSH
  call MPI_INIT
  call MPI_COMM_RANK ( MPI_COMM_WORLD,
    myrank)
  call MPI_COMM_GROUP( MPI_COMM_WORLD,
    mpi_group_world)
  read nprocs_PIMD(m) from input file
  read nprocs_SMASH(n) from input file

! Set PIMD communicator
  do i=1, nprocs_PIMD
     ranklist_PIMD(i)= nprocs_SMASH*(i-1) + mod(myrank,nprocs_SMASH)
  enddo
  call MPI_GROUP_INCL ( mpi_group_world, nprocs_PIMD,
   ranklist_PIMD, mpi_group_PIMD)
  call MPI_COMM_CREATE ( MPI_COMM_WORLD, mpi_group_PIMD, mpi_comm_PIMD)
  call MPI_COMM_RANK ( mpi_comm_PIMD, myrank_PIMD)

! Set SMASH communicator
  do i=1, nprocs_SMASH
    ranklist_SMASH(i)= int(myrank/nprocs_SMASH)*nprocs_SMASH + i - 1
  enddo
  call MPI_GROUP_INCL ( mpi_group_world, nprocs_SMASH,
   ranklist_SMASH, mpi_group_SMASH)
  call MPI_COMM_CREATE ( MPI_COMM_WORLD, mpi_group_SMASH, mpi_comm_SMASH)
  call MPI_COMM_RANK ( mpi_comm_SMASH, myrank_SMASH)

! Start PIMD-SMASH calculation
  do istep=1, maxstep
    do ibead=1, nbead
      if( mod(ibead-1,nprocs_PIMD).ne.myrank_PIMD) cycle
      call PIMD (simulation_data[ibead](inout), geometry[ibead](inout),
        energy[ibead](in), force[ibead](in))
      call SMASH (mpi_comm_SMASH(in), geometry[ibead](in),
        energy[ibead](out), force[ibead](out))
    enddo
    call MPI_ALLREDUCE (mpi_comm_PIMD(in),energy(inout))
    call MPI_ALLREDUCE (mpi_comm_PIMD(in),force(inout))
  enddo
end program
```

c)

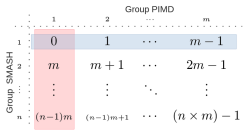

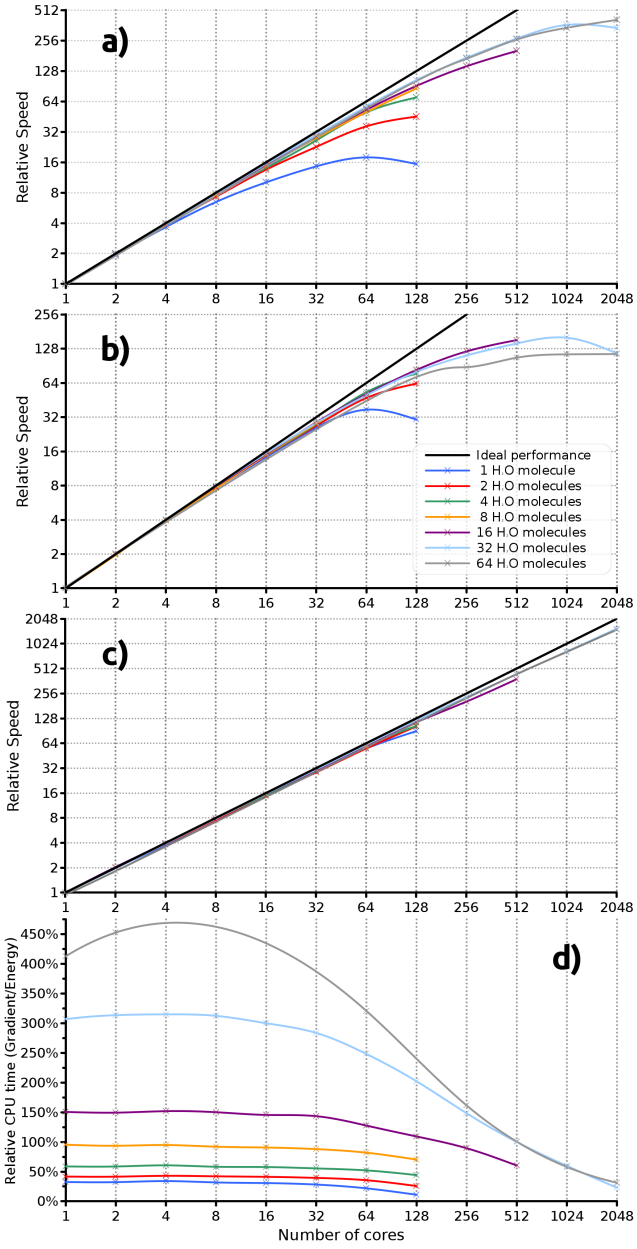

Figure 2: The parallel performance of path integral MD simulations of 1–64 water molecules with two beads. (a) The relative speed in the total simulation. The relative speed in the calculation of (b) the BO energy and (c) the BO force. (d) The ratio of the average time spent on the BO force and the BO energy.

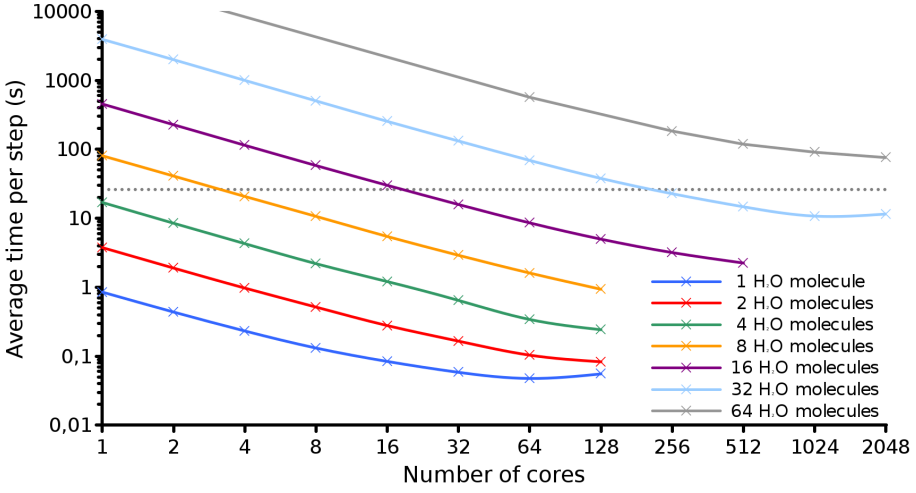

Figure 3: Average time per step in path integral MD simulations of 1–64 water molecules with two beads. The colour codes are the same as those of Figure 2.

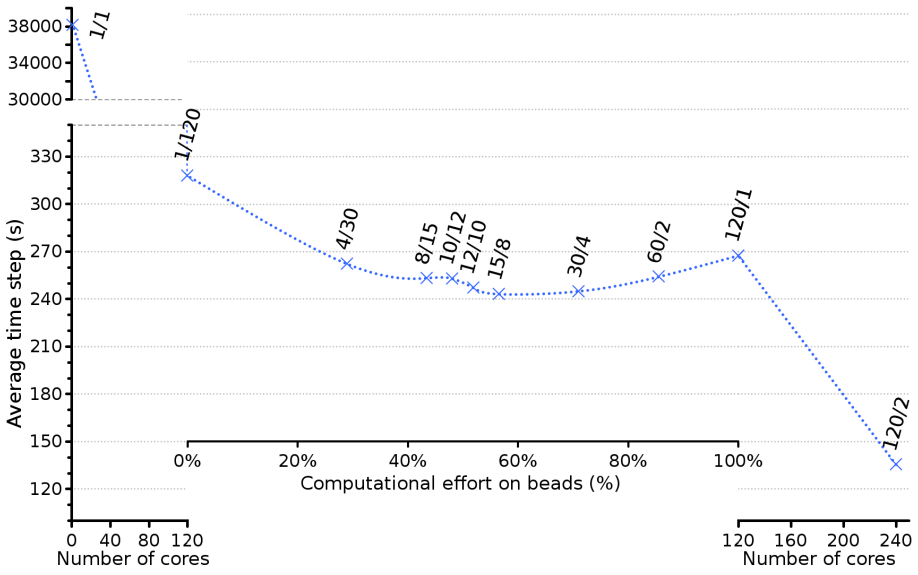

Figure 4: Average time per step of the string method for a DA reaction of cyclopentadiene and butenone. 120 images are used to describe the structural change along the string. The image parallel number n and the force-parallel number m are given as "n/m" where $n \times m$ is the total number of cores used. The index $\log_{120}(n)$ measures the computational effort put on the image parallel number over the force-parallel number for a fixed number of cores $n \times m = 120$.

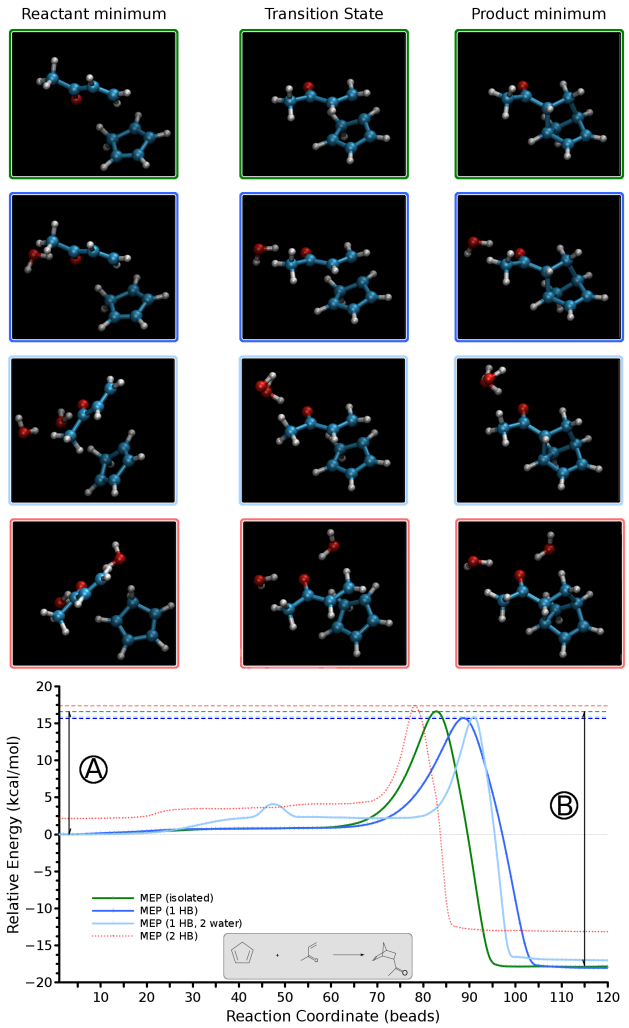

Figure 5: The minimum energy paths of the DA reaction of cyclopentadiene and butenone obtained by the string method. The results are shown for the cases without water (green), with one water molecule (dark blue) and with two water molecules in a chain (light blue). The path for the case of with two water forming a double hydrogen bond with the ketone in plotted with a dotted red line. Zero energy is set to that of the reactant minimum of each molecular cluster. The respective molecular structures of the reactant, transition, and product states are also given. The activation energy of forward reaction (Ⓐ): 16.6 kcal/mol (without water), 15.7 kcal/mol (with one water molecule) and 15.8 kcal/mol (with two water molecules). The activation energy of backward reaction (Ⓑ): 34.5 kcal/mol (without water), 33.7 kcal/mol (with one water molecule) and 32.8 kcal/mol (with two water molecules).